Nonlinear Computing with Optical Twinning Trees

Denys I. Bondar[1] and Gerard McCaul[2]
[1] Department of Physics and Engineering Physics, Tulane University, New Orleans, Louisiana 70118, USA
[2] EigenEngine

Can two different materials be made to look optically identical under exactly the same driving field? A twinning field does so for a chosen response, but in the original formulation it is obtained by calculating the drive in advance from models of both materials. We show that the field can instead be found by a causal optical loop. A balanced detector measures the difference between the two responses; a signed integral controller returns that difference to the shared drive. When the loop settles, the measured difference vanishes and the retained drive is the twinning field. The loop is locally stable for $\kappa|\chi_\Delta|\tau < \pi/2$. Nonlinear sweeps recover this boundary, and exact propagation of two four-site Hubbard rings shows their currents locking under one common Peierls phase. The same idea can also be run backwards: choose the materials so that their twinning fields are the answers one wants. We select a Kerr and a linear cavity whose equal-intensity points encode the non-negative roots of a supplied quadratic. Two such conditions sharing one reference cavity form a product residual—a twinning tree—whose zeros are the roots of a supplied factorised quartic. The full three-cavity, seven-variable dynamics acquire each of its four roots; 33 of 42 initial-condition and polarity trials reach a root, while the remainder meet imposed power bounds. Coalescing roots erase the differential susceptibility, making even-multiplicity roots one-sided under fixed-polarity feedback. These results are numerical and analytic; the optical circuits are proposed implementations.

I TWINNING FIELDS IN NONLINEAR OPTICS

A material is often introduced optically by a list of things it does to light. That shorthand is powerful because the weak-field answer is reproducible. Under strong driving, however, the interrogating field becomes a participant: it can redirect a many-body current [1], force a chosen observable trajectory [2], or suppress one component of a composite response [3]. The map from field to nonlinear response need not even be one-to-one [4]. At that point the familiar division between a property of the material and a property of the experiment begins to soften. If the field helps to make the answer, how much of a material's optical identity belongs to the material alone?

Twinning fields give one answer. A twinning field is a single electromagnetic drive $\Phi_\tau(t)$ for which two distinct systems return the same selected response,

$$J_1[\Phi_\tau](t) = J_2[\Phi_\tau](t).$$

The construction was introduced for nonlinear many-body systems and demonstrated for Fermi-Hubbard and tight-binding materials [5]. Its insistence on one identical field is the point. With two independently tailored fields, tracking control can make different systems reproduce the same target by correcting each one separately [1, 2]. A twinning field instead identifies a trajectory on which two distinct response laws coincide under one shared physical cause.

In the original formulation, that trajectory is calculated in open loop from the evolving states of both systems [5]. This establishes the object, but leaves an experimental loop to close. The field must be known before it is applied, while the information used to calculate it is returned by the responses produced after application. Feedback resolves the apparent circle by allowing the materials to supply that information as the experiment proceeds. Tracking feedback can generate a prescribed broadband optical behaviour [6], and a measured response mismatch can be amplified and returned without first shaping a correction pulse [7]. Those architectures regulate one response or correct one branch towards another. The stricter task here is to adapt the one field incident on both systems.

Once the problem is written this way, a second interpretation becomes difficult to avoid. Balanced detection measures $J_1 - J_2$. Laser locking calls such a signed difference an error signal [8, 9]; numerical analysis calls it a residual. A controller that adjusts Φ until $J_1 - J_2 = 0$ is therefore finding a root, with the two physical systems evaluating the function. The shift is physical as well as linguistic: the field that interrogates the materials becomes the unknown, and their disagreement supplies each iteration. There is already a substantial optical-computing world around this observation. Designed electromagnetic media perform mathematical operators [10] and solve integral equations [11]; coherent feedback supplies analogue solutions of partial differential equations [12], microring networks solve ordinary differential equations [13], and inverse-designed free-space feedback solves integral equations [14]. Delayed and integrated photonic reservoirs process information in their dynamics [15, 16], while optoelectronic recurrence searches for computational fixed points [17, 18]. The question here is consequently narrow: can equality between *distinct common-driven responses* encode the equation, and can the same shared field acquire its solutions?

We begin by treating a twinning field as the zero of a general differential response. Linearising the loop exposes a gain-delay boundary, after which exact propagation of two four-site Hubbard rings tests common-field acquisition in the many-body setting that motivated the proposal. Having established that the field can be found, we ask whether it can find anything in return. Nonlinear optical response has previously been recruited as a computational transformation [19], and Hamiltonian parameters can encode data for a physical processor [20]. Here the coefficients of a supplied quadratic choose the input-output laws of a Kerr and a linear cavity. Their nonzero equal-intensity points occur at its non-negative roots. The controller receives only the two measured intensities: the equation is evaluated by the optical responses, not by a digital call to the polynomial.

The same construction admits a higher-order extension. Two Kerr cavities can share one linear reference, with each equality condition representing one supplied quadratic factor. Multiplying their normalised response differences produces a residual that vanishes at the union of the two root sets. Layered optics can already generate higher polynomial orders [21], analogue optical multiplication is established [22], and programmable photonic nonlinearities can represent coefficient-dependent polynomial transformations [23]. Organising these elements around material twinning within one common-field loop produces a seven-variable calculation in which three complex cavity amplitudes and one real controller state acquire all four roots of a supplied factorised quartic. Pairwise optical indistinguishability has become a composable physical calculation. The factorisation is part of the encoding. Arbitrary expanded-polynomial encoding, automatic factorisation, scaling with hierarchy depth, and optical advantage are not established.

The distinction between what is derived, computed, and merely proposed will be maintained throughout. Section II defines strict common-field acquisition, Sec. III treats the effect of delay, and Sec. IV returns to twinned Hubbard materials. Section V constructs and tests the quadratic cavity encoder. Section VI composes two such conditions into the factorised-quartic calculation, and Sec. VII examines the loss of directional information when roots coalesce. Section VIII places these results among neighbouring approaches to optical computation and identifies the physical questions left by the construction.

II A TWINNING FIELD AS A PHYSICAL ROOT

Consider two systems with internal states $x_1(t)$ and $x_2(t)$. Their response to a drive depends both on its present value and on the state prepared by its history, which we write as $J_j[\Phi(t), x_j(t)]$. We define strict twinning by requiring the same total field $\Phi(t)$ to reach both systems and the chosen responses to agree over a declared observation interval $\mathcal{T}_{\text{obs}}$,

$$J_1[\Phi(t), x_1(t)] = J_2[\Phi(t), x_2(t)], \qquad t \in \mathcal{T}_{\text{obs}}.$$

This is stronger than asking one branch to imitate the other using an additional correction field. The equality concerns what the two systems do under one shared physical cause [5].

The quantity available to a balanced detector is the differential response

$$e(t) = J_1[\Phi(t), x_1(t)] - J_2[\Phi(t), x_2(t)]. \tag{1}$$

Equation (1) is simultaneously the optical error signal and the residual of the root problem. This identification carries more than a change of vocabulary. The unknown Φ is not passed to a numerical routine together with analytic functions J_1 and J_2. It is applied to the systems, and their measured difference evaluates e. Signed optical zero crossings are the ordinary currency of cavity locking [8, 9]; the change here is that the two arms contain distinct response evaluators and their equality, rather than one cavity resonance, defines the lock point.

There is nevertheless a causal price. At time t, the controller has access only to measurements that have returned by t and to its own stored state. With a total propagation, detection, and processing delay $\tau > 0$, it cannot know the response to the field it is applying at that same instant. Equality from an arbitrary initial time is therefore unavailable to a causal feedback loop. What can be done is to acquire a root over an interval $\mathcal{T}_{\text{acq}}$ and then test the resulting field over $\mathcal{T}_{\text{obs}}$. Delayed response information is a standard ingredient of nonlinear control [24] and of computational feedback dynamics [15]; related optical response-control schemes use it directly [6, 7]. The distinction here is that the update acts before the passive split, so both systems continue to receive exactly the same field.

The simplest controller that retains a nonzero field when the error vanishes is an integrator,

$$\dot{\Phi}(t) = -\kappa\sigma e(t - \tau), \tag{2}$$

where $\kappa > 0$ is the gain and $\sigma = \pm 1$ fixes the feedback polarity. A proportional correction would disappear at $e = 0$ and return the apparatus to its uncorrected drive. By contrast, Eq. (2) stores the acquired value in the controller state. At an interior stationary point,

$$\dot{\Phi} = 0 \quad \Longleftrightarrow \quad e(\Phi_\star) = 0,$$

provided the two systems have settled into the corresponding driven states. A regular twinning field is thus a fixed point of the loop, and a stable fixed point is a twinning field for the selected response [2].

This equivalence also fixes what we mean by success. During $\mathcal{T}_{\text{acq}}$, returned measurements are permitted to change Φ. Once a residual threshold has been maintained for a prescribed hold time, the controller state is retained and the equal-response claim is assessed on the subsequent observation interval. A transient crossing of $e = 0$ is not a solution, and equality during acquisition is not asserted before information has had time to circulate.

Figure 1 puts this distinction into an optical layout. A common modulator prepares Φ, a passive split sends it to the two systems, and balanced detection forms Eq. (1). The signed integrator implements Eq. (2). Every component has a close optical precedent in cavity locking and feedback computing [9, 14]; the location of the split is what preserves the strict common drive. Open-loop synthesis, asymmetric correction, and strict common-field acquisition therefore differ in where the controllable field is introduced, not merely in the electronics used to calculate it. The remaining question is whether the root survives the delay needed to observe it.

III DELAY IN THE COMMON-FIELD LOOP

Feedback succeeds only if the measured difference changes in a useful direction when the field is perturbed. Let $\Phi_\star$ be a twinning root and freeze the internal-state dependence over the local stability calculation. The differential susceptibility is

$$\chi_\Delta = \partial_\Phi [J_1 - J_2]|_{\Phi_\star} . \tag{3}$$

Equation (3) is the slope of the physical residual. Its sign selects the restoring polarity, while its magnitude tells the controller how much information about displacement from the root is present in one differential measurement.

We write $\Phi(t) = \Phi_\star + \delta\Phi(t)$ and expand Eq. (1) to obtain

$$e(t) = \chi_\Delta\, \delta\Phi(t) + O(\delta\Phi^2).$$

With $\sigma\chi_\Delta > 0$, the controller in Eq. (2) therefore reduces locally to

$$\delta\dot{\Phi}(t) = -\kappa|\chi_\Delta|\, \delta\Phi(t-\tau).$$

At zero delay, the root contracts exponentially at rate $\kappa|\chi_\Delta|$. Increasing the gain speeds acquisition only until the delayed correction begins to act on obsolete information.

Taking $\delta\Phi \propto e^{\lambda t}$ yields the characteristic equation and its stable branch [25],

$$\lambda + \kappa|\chi_\Delta| e^{-\lambda\tau} = 0, \qquad \kappa|\chi_\Delta|\tau < \frac{\pi}{2}. \tag{4}$$

The boundary follows directly by setting $\lambda = i\omega$. Separating real and imaginary parts gives

$$\cos(\omega\tau) = 0, \qquad \omega = \kappa|\chi_\Delta| \sin(\omega\tau),$$

so the first crossing occurs at $\omega\tau = \pi/2$ and $\omega = \kappa|\chi_\Delta|$. Appendix A records the contraction and fitting conventions used below.

The result in Eq. (4) is local and scalar. It does not assume that the full nonlinear dynamics remain frozen far from the root, and it does not guarantee acquisition within a finite experimental window. Delayed feedback can stabilise or destabilise nonlinear dynamics depending on precisely this competition between response, gain, and latency [15, 24]. Equation (4) identifies the dimensionless combination against which the numerical calculations should be organised.

We test the boundary with a 134-point nonlinear delayed sweep. The fitted residual exponent changes sign at $\kappa|\chi_\Delta|\tau = 1.5675$, within 0.0033 of $\pi/2$. A practical finite-time acquisition test ceases near 1.53, where convergence is still asymptotically stable but too slow to satisfy the allotted hold time. Figure 2 shows both criteria and representative trajectories below, near, and above the crossing. Delay does not alter what the root is. It determines whether the apparatus can remain there.

IV TWINNING TWO HUBBARD MATERIALS

The scalar calculation establishes how a regular differential-response root behaves, but it says little about where that response comes from. We therefore return to the many-body setting of the original twinning field problem. The two response evaluators are half-filled four-site spinful Fermi-Hubbard rings,

$$H_j(\Phi) = -t_0 \sum_{\ell,\varsigma} \left(e^{-i\Phi} c^\dagger_{\ell\varsigma} c_{\ell+1,\varsigma} + \text{H.c.} \right) + U_j \sum_\ell n_{\ell\uparrow} n_{\ell\downarrow}. \tag{5}$$

Equation (5) is evaluated with $U_1 = 0.5$ and $U_2 = 2.0$, so the systems are dynamically distinct but share the same lattice and hopping scale. The common Peierls phase $\Phi(t)$ is the field to be acquired. This response channel descends from driven-system tracking and current control [1, 2]; the use of equality under one field is inherited from the twinning-field construction [5].

For each ring, the current is the response conjugate to the phase,

$$\hat{J}_j(\Phi) = -\partial_\Phi H_j(\Phi), \qquad J_j(t) = \langle\psi_j(t)|\hat{J}_j[\Phi(t)]|\psi_j(t)\rangle,$$

and the states evolve independently according to

$$i\partial_t|\psi_j(t)\rangle = H_j[\Phi(t)]|\psi_j(t)\rangle.$$

The only coupling between the two calculations is the shared phase updated from $J_1 - J_2$. Fixed particle number and spin reduce each ring to a 36-state sector. This small system is large enough for the interactions to produce different nonlinear currents, while remaining tractable under exact propagation.

At each time step, the delayed current difference is inserted into Eq. (2). The controller does not receive either wavefunction, the value of U_j, or an independently calculated twinning field. For comparison only, we separately invert the instantaneous current-equality condition using the propagated states. This produces a reference phase $\Phi_{\text{inst}}(t)$ against which the acquired phase can be tested without feeding the answer back into the loop. State

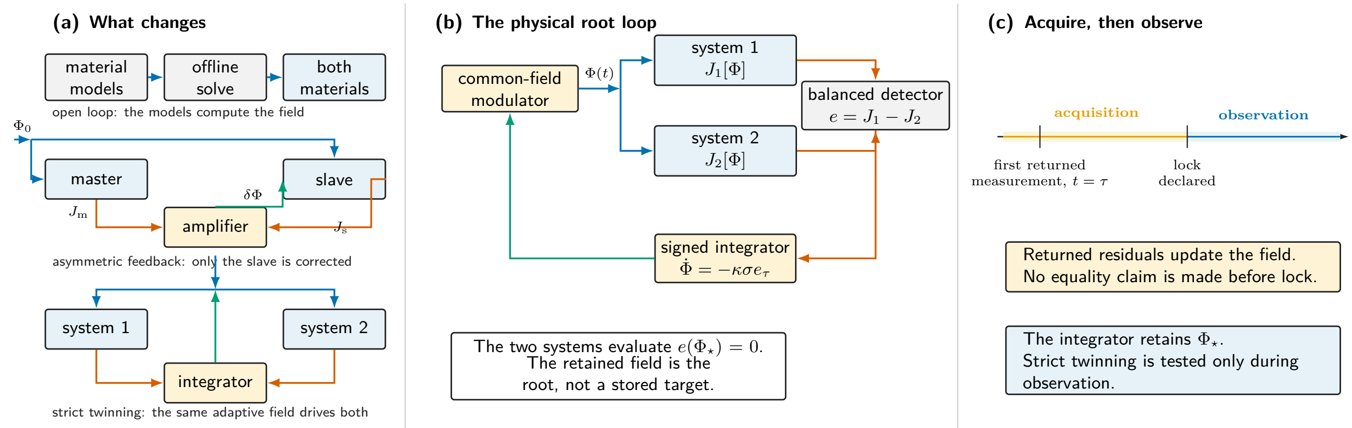

FIG. 1. The common field feedback architecture illustrates how a twinning field becomes a physical root. (a) Open-loop synthesis solves for a field in advance, asymmetric master–slave feedback corrects one branch, and strict twinning instead sends one identical total field to both systems. (b) The two systems evaluate the differential response $e = J_1 - J_2$, a balanced detector measures it, and a signed integrator updates the shared field. (c) Nonzero latency introduces an acquisition interval before the observation window on which strict twinning is asserted. The diagram specifies a proposed architecture rather than an experimental realisation.

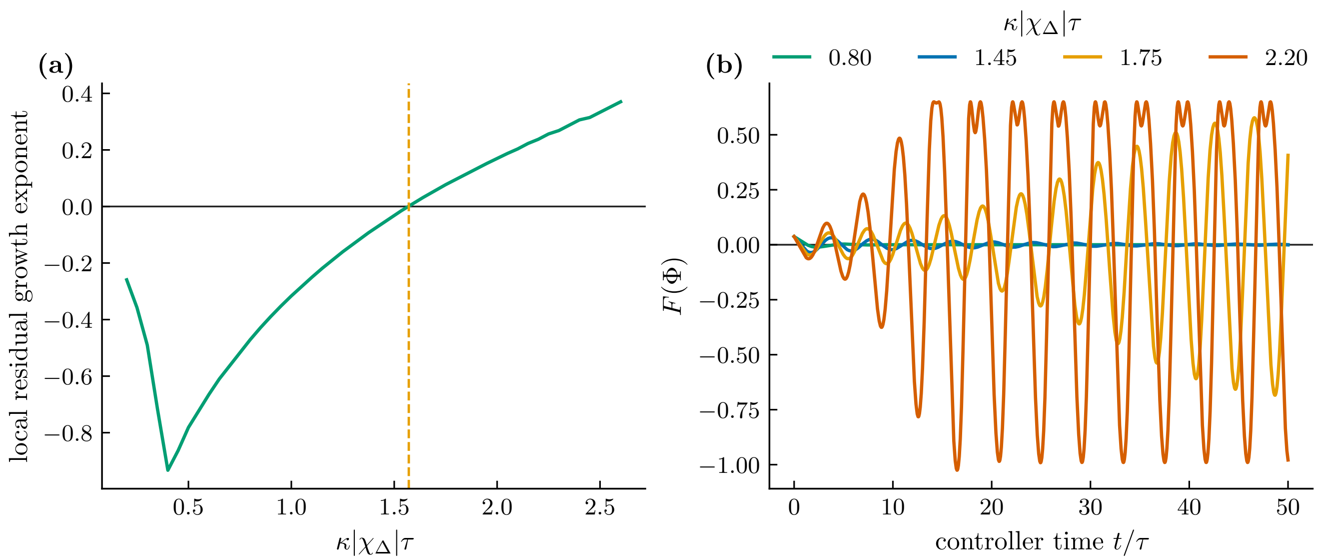

FIG. 2. The delayed scalar calculation shows the gain–delay stability boundary. (a) The fitted residual exponent changes sign at 1.5675, within 0.0033 of the analytic value $\pi/2$. The more conservative finite-time acquisition criterion ends near 1.53. (b) Representative traces show persistent acquisition below the boundary, slow decay near it, and growing oscillations above it. All curves are projected from retained numerical tables.

preparation, integration, and error measures are given in Appendix B.

The two currents lock over the observation interval. Their tail root-mean-square mismatch is 0.50% of the root-mean-square common current, and the acquired phase follows $\Phi_{\mathrm{inst}}(t)$ with root-mean-square error 0.00732. These tests answer different questions. The first establishes optical equality. The second checks the interpretation developed in Sec. II: the controller state tracks the root that the two evolving materials define.

Figure 3 shows the currents, the differential response, the acquired and independently inverted phases, and the loss of practical lock in a gain sweep. The initial susceptibility-scaled gain-delay value at loss of lock is approximately 1.64. Its proximity to the scalar $\pi/2$ boundary in Eq. (4) is striking given that the states continue to evolve while the field is acquired. It is nevertheless a numerical correspondence, not a stability theorem for the Hubbard dynamics. The many-body calculation shows that a twinning root can be acquired. We next arrange

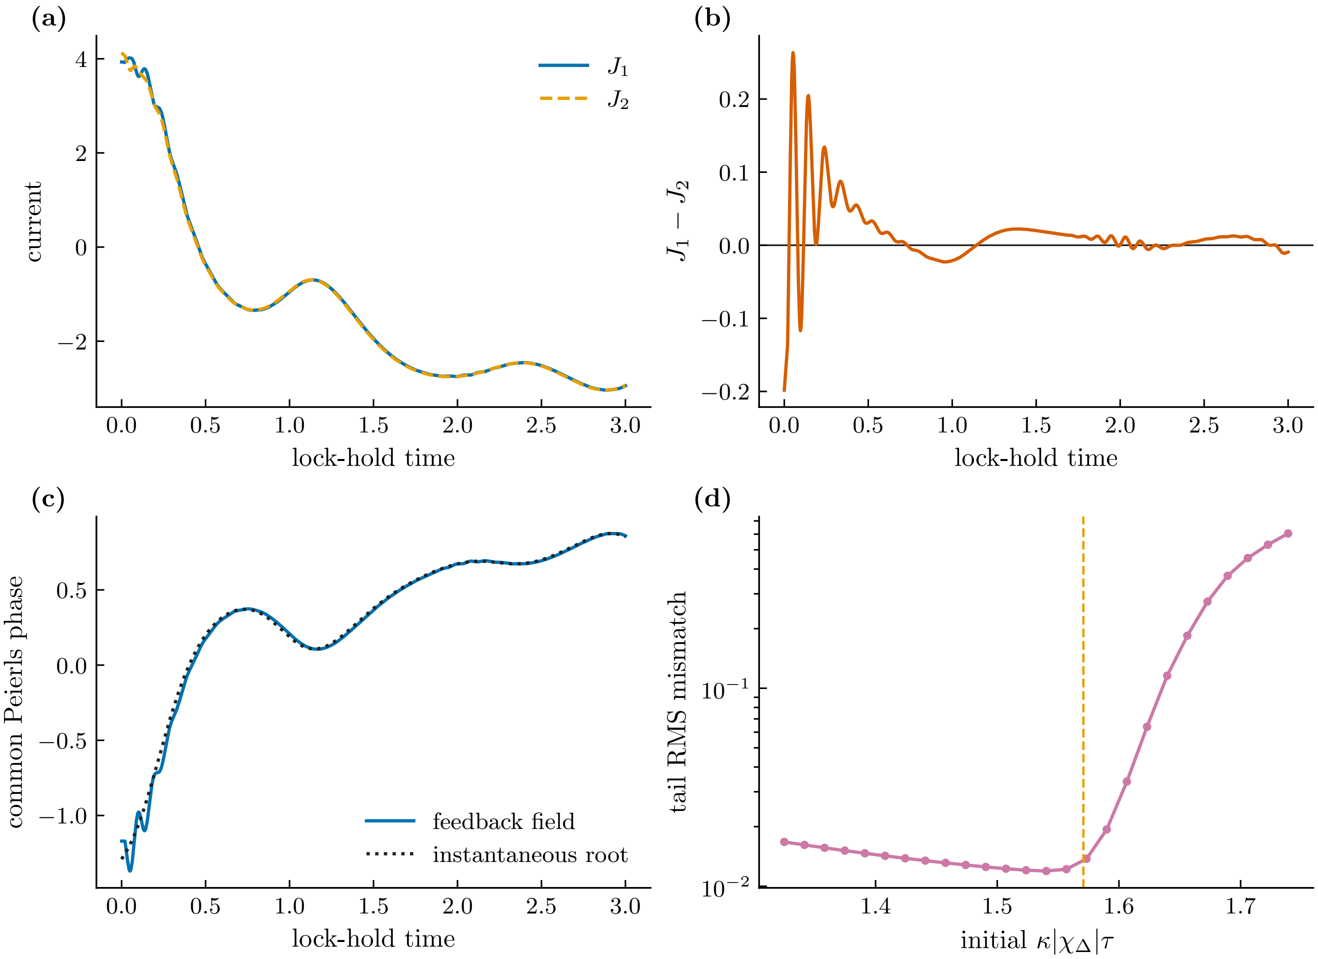

FIG. 3. The four-site Hubbard calculation shows acquisition of a common twinning field. Two half-filled spinful rings, each represented in a 36-state symmetry sector and with interactions $U_1 = 0.5$ and $U_2 = 2$, receive the same Peierls phase. (a) Their currents become indistinguishable on the plotted scale. (b) The difference decays to a tail RMS mismatch equal to 0.50% of the common-current RMS. (c) The acquired phase follows the independently inverted instantaneous twinning root with RMS error 0.00732. (d) A gain sweep loses practical lock above an initial susceptibility-scaled gain–delay product of approximately 1.64, close to the scalar boundary.

for the root itself to carry a prescribed equation.

V A QUADRATIC ENCODED BY TWO CAVITIES

Suppose the desired roots are those of a supplied real quadratic

$$q(n) = n^2 + bn + c.$$

A direct numerical calculation would evaluate q and adjust n. The optical construction instead uses two distinct response laws whose equality is equivalent to $q(n) = 0$. Kerr cavities are a natural candidate. Optical bistability was observed in a nonlinear Fabry–Pérot resonator [26], given a quantum input–output model [27], and shown to support multivalued stationary states in a ring cavity [28]; the Lugiato–Lefever equation established the wider driven-dissipative Kerr-cavity setting [29]. The extra order will shortly prove useful rather than troublesome.

For a coherently driven single-mode cavity, the mean field satisfies

$$\begin{aligned} \dot{a}_j &= \left[-\gamma_j + i\left(\Delta_j - K_j|a_j|^2\right)\right] a_j + \eta, \\ P_j(n) &= n\left[\gamma_j^2 + (\Delta_j - K_j n)^2\right] = |\eta|^2, \end{aligned} \tag{6}$$

where $n = |a_j|^2$, γ_j is the linewidth, Δ_j is the detuning, and K_j is the Kerr coefficient. Equation (6) makes $P_j(n)$ the input power required to sustain intensity n on a steady branch. Two cavities driven by the same η have equal intensities wherever their functions $P_1(n)$ and $P_2(n)$ coincide.

Take cavity 1 to be Kerr nonlinear with $K_1 = 1$, and cavity 2 to be linear with $K_2 = 0$ and $\Delta_2 = 0$. Their

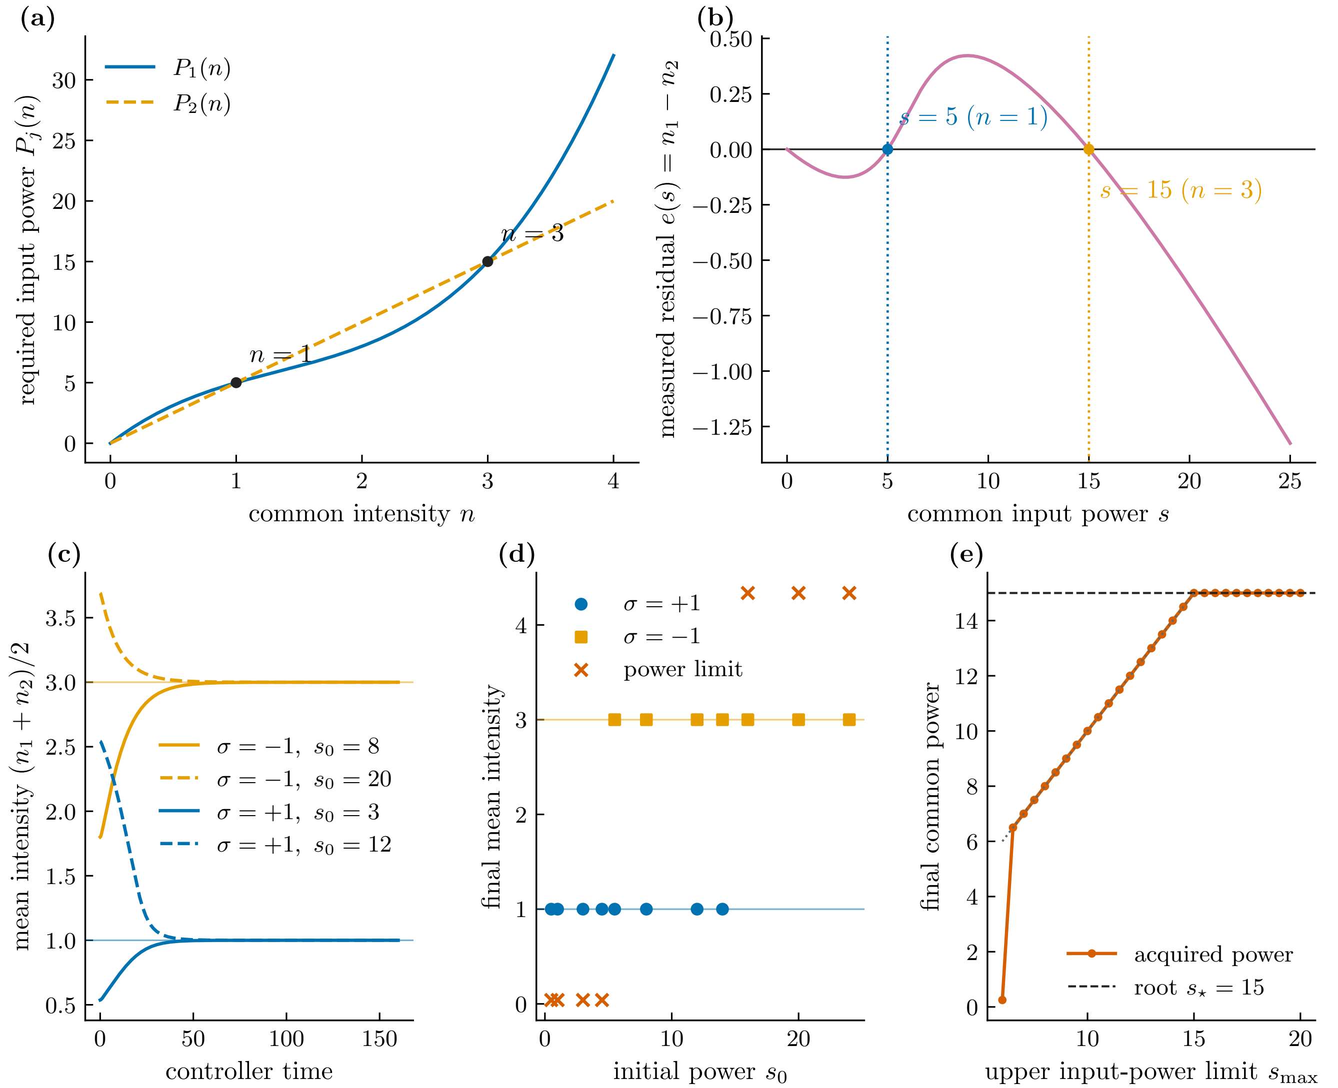

FIG. 4. The common-driven cavity pair solves the encoded quadratic $n^2-4n+3=0$. (a) The cavity input–output laws intersect at $n = 1$ and 3. (b) Their measured intensity difference vanishes at common powers $s = 5$ and 15. (c) Representative full complex-amplitude trajectories acquire $n = 1$ for positive polarity and $n = 3$ for negative polarity. (d) Fifteen of 22 multi-start runs acquire one of those roots with maximum intensity error 2.18×10^{-6}, while wrong-basin trajectories reach an imposed input power limit. (e) Seventeen of 18 upper limits below $s = 15$ exclude the high-power root, whereas all ten limits above it permit acquisition.

input-output difference is

$$P_1(n) - P_2(n) = n\left[(\gamma_1^2 + \Delta_1^2 - \gamma_2^2) - 2\Delta_1 n + n^2\right].$$

Matching the bracket to $q(n)$ fixes the detuning and linewidth difference. One convenient parameterisation is

$$\begin{aligned} K_1 = 1, \quad K_2 = 0, \qquad \Delta_1 = -b/2, \quad \Delta_2 = 0, \\ \gamma_1^2 = L, \qquad \gamma_2^2 = L + \Delta_1^2 - c, \\ P_1(n) - P_2(n) = nq(n). \end{aligned} \tag{7}$$

The free positive offset L keeps both linewidths physical and can be chosen so that the working branches remain monotone.

Equation (7) has an unavoidable zero at $n = 0$, where two empty cavities are trivially indistinguishable. Away from this vacuum solution,

$$P_1(n) = P_2(n) \iff q(n) = 0, \qquad n > 0.$$

Thus each real non-negative root r is represented by an equal-intensity state of the two cavities. The corresponding control variable is not r itself but the common input power

$$s_\star = P_1(r) = P_2(r).$$

The coefficients determine the cavity parameters, the cavity responses evaluate the residual, and the field value retained by the controller identifies the selected solution. Hamiltonian-parameter encoding and programmable optical nonlinearities provide broader precedents for choosing a physical response law from supplied data [20, 23]; the equality-selected root is the present construction.

During acquisition, the full complex amplitudes are evolved rather than replaced by their steady curves. The intensity difference is the cavity realisation of the residual in Eq. (1). If $n_j(s,t)$ denotes the returned intensity under common power s, the controller is

$$\dot{s} = -g\sigma[n_1(s,t) - n_2(s,t)], \qquad \delta\dot{s} = -g\sigma e'(s_\star)\delta s. \quad (8)$$

Equation (8) also gives the local linearisation. Simple roots on opposite sides of a quadratic generally have opposite slopes and therefore use opposite restoring polarities. The polarity supplies a basin choice, not the root value.

For $q(n) = n^2 - 4n + 3 = (n-1)(n-3)$, the map gives $P_1 - P_2 = n(n-1)(n-3)$. The two nontrivial roots correspond to common powers $s = 5$ and $s = 15$. Full cavity dynamics acquire $n = 1$ with one polarity and $n = 3$ with the other, as predicted by Eq. (8). Across 22 initial-power and polarity combinations, 15 acquire a root with maximum intensity error 2.18×10^{-6}. The remaining trajectories are launched in the wrong basin and reach an imposed power limit. They do not converge to spurious equalities.

Figure 4 connects the algebra to the dynamics. It shows the encoded polynomial, the physically measured differential intensity, representative acquisitions, and the polarity-resolved basins. An additional power-limit sweep finds that 17 of 18 upper limits below $s = 15$ exclude the higher root, while all ten limits above it permit acquisition. The optical loop therefore solves only within the range it can physically explore. This differs from material solvers in which a metastructure performs an operator or equation [10, 11], from inverse-designed free-space networks for integral equations [14], and from feedback resonators for differential equations [12, 13]. It also differs from delay-based photonic reservoirs trained for information processing [15, 16], and from recurrent optoelectronic or analogue optical machines seeking optimisation fixed points [17, 18]. The present residual is the difference between two material response laws. Appendix B records the parameter bounds, integration details, and retained tables.

VI OPTICAL TWINNING TREES

The cavity pair in Sec. V represents a quadratic because the difference of its two input-output laws contains one quadratic factor. The same Kerr-linear pair cannot represent a quartic response difference. There is another route. If two independently measured residuals depend on the same control variable, their product vanishes wherever either residual vanishes,

$$\mathcal{Z}(e_A e_B) = \mathcal{Z}(e_A) \cup \mathcal{Z}(e_B),$$

where $\mathcal{Z}$ denotes the set of zeros. The order of the encoded equation can therefore be increased by composing response equalities rather than increasing the order of any one material response.

To make this economical, both quadratic conditions share a linear reference cavity with input-output law $P_0(n) = Gn$. Let

$$q_j(s) = s^2 + b_j s + c_j$$

be a supplied factor written in terms of the common input power s. Choosing the detuning and linewidth of Kerr cavity j as

$$\begin{aligned} \Delta_j &= -\frac{b_j K_j}{2G}, \\ \gamma_j^2 &= G - \Delta_j^2 + \frac{c_j K_j^2}{G^2}, \\ P_j(n) - P_0(n) &= \frac{n K_j^2}{G^2} q_j(Gn) \end{aligned} \quad (9)$$

makes the equality $P_j = P_0$ occur at the non-negative roots of $q_j(s)$. Equation (9) is the shared-reference counterpart of the coefficient map in Eq. (7). It follows by substituting the Kerr input-output law of Eq. (6) and matching powers of Gn, and it uses the supplied factor coefficients directly.

One coherent input now drives a reference cavity and two Kerr cavities A and B. Balanced detection forms their differences from the reference, after which a multiplier and one outer integrator produce

$$\begin{aligned} e_A &= |a_A|^2 - |a_0|^2, \quad & e_B &= |a_B|^2 - |a_0|^2, \\ E &= \frac{e_A}{\epsilon_A}\frac{e_B}{\epsilon_B}, & \dot{s} &= -g\sigma E. \end{aligned} \quad (10)$$

The scales ϵ_A and ϵ_B make the two detector outputs dimensionless and prevent one leaf from dominating solely by units. Analogue optical multiplication has a component-level precedent [22], while layered optical systems already use repeated multiplication to build polynomial order [21]. Figure 5 shows how those operations are arranged around one shared reference and one control field.

For the numerical example we choose $G = 5$ and

$$\begin{aligned} q_A(s) &= s^2 - 3s + 2 = (s-1)(s-2), \\ q_B(s) &= s^2 - 7s + 12 = (s-3)(s-4). \end{aligned}$$

Equations (9) and (10) then give

$$\begin{aligned} E(s) = 0 \quad &\Longleftrightarrow \quad Q(s) = 0, \\ Q(s) &= q_A(s) q_B(s) \\ &= (s-1)(s-2)(s-3)(s-4), \end{aligned}$$

within the monotone operating branches and excluding the common vacuum equality. The dynamical model contains three complex amplitudes and the real controller

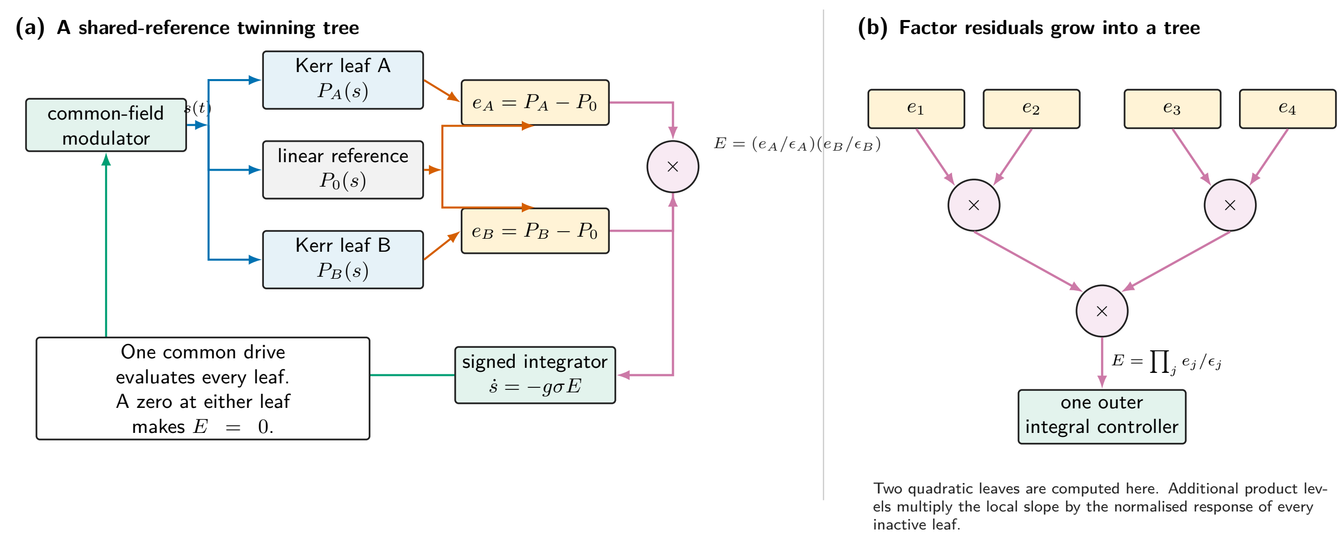

FIG. 5. The shared-reference construction illustrates how several twinning conditions feed one common field controller. (a) One coherent input reaches a linear reference and two distinct Kerr cavities. Balanced detection forms $e_A = P_A - P_0$ and $e_B = P_B - P_0$, an optical or electronic multiplier forms $E = (e_A/\epsilon_A)(e_B/\epsilon_B)$, and one integral controller updates the common input. The product E vanishes when either leaf response equals the reference. After this construction is established, we refer to the hierarchy as a twinning tree. (b) Further product stages state the recursive construction rule. The present calculation demonstrates two leaves and does not establish scaling with depth.

state s, giving seven real variables. It neither discovers the factorisation nor evaluates an arbitrary expanded quartic.

All four roots are acquired in the full three-cavity dynamics. Of 42 initial-power and polarity combinations, 33 acquire a root, four reach the lower imposed power limit, five reach the upper limit, and none remain unlocked. The largest acquired power error is 9.65×10^{-4} at the slowest-conditioned root. Repeating representative trajectories with time steps 0.01, 0.005, and 0.0025 leaves every classification unchanged, with maximum final-power shift 3.16×10^{-12} between the two finest steps. Figure 6 shows the two leaf residuals, their four-zero product, the acquired trajectories, and the polarity-resolved basins. Appendix B gives the parameter map and numerical custody.

We have now earned a compact name for the construction. A hierarchy of material response-equality conditions, joined by shared references and product residuals under one optical control, will be called a *twinning tree*. The term does not rename its components. Numerical analysis still sees a product residual, and analogue photonics still sees a multiplier tree. The twinning tree is their organisation around the equality of distinct driven responses. Figure 5 sketches a recursive extension, while Fig. 6 demonstrates only the two-leaf case. Before adding more leaves, it is worth asking how reliably even these four roots can be found.

VII WHEN OPTICAL ROOTS BECOME ILL-CONDITIONED

The feedback loop learns which side of a root it occupies from the sign of the differential response. This information weakens when two roots approach each other. For a quadratic $q(n) = (n - r_-)(n - r_+)$, the slopes at the roots are

$$q'(r_\pm) = \pm(r_+ - r_-).$$

The sign still distinguishes the two roots, but the magnitude of the distinction collapses with their separation.

For the cavity residual, implicit differentiation of $P_j[n_j(s)] = s$ gives $\partial_s n_j = 1/P_j'(n_j)$. Applying this to the intensity difference at a simple root yields

$$\chi_\Delta^{(\pm)} = -\frac{r_\pm q'(r_\pm)}{P_1'(r_\pm)P_2'(r_\pm)}, \qquad q'(r_\pm) = \pm(r_+ - r_-). \quad (11)$$

Equation (11) connects the algebraic conditioning of the polynomial to the physical susceptibility defined in Eq. (3). Provided the two cavity branches remain regular, the local feedback rate vanishes linearly as the roots coalesce. The regular-branch qualification matters because nonlinear cavities can be multivalued [26, 28] and dynamically unstable [27] independently of the encoded polynomial.

The numerical consequence is direct. All seven tested simple root cases settle, but the measured settling time grows from 36.09 to 210.14 over the sampled separations. At exact even multiplicity, the residual has the same sign on both sides of the root,

$$q(n) = (n - r)^2 \geq 0.$$

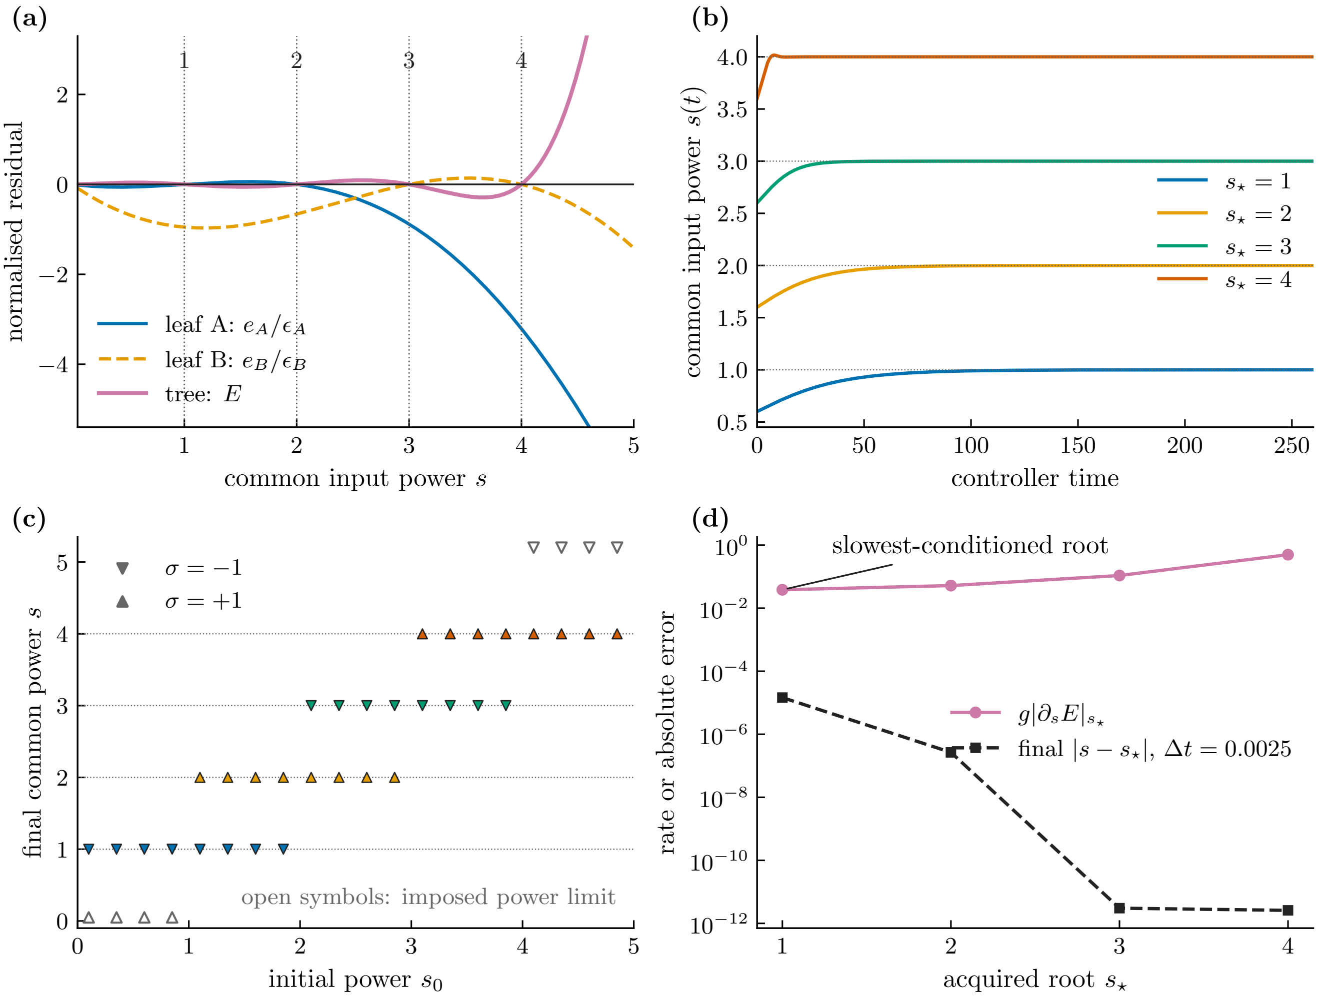

FIG. 6. The seven-variable shared-reference calculation acquires the four roots of $(s-1)(s-2)(s-3)(s-4) = 0$. (a) The product vanishes where either differential response vanishes. Leaf A supplies $s = 1, 2$ and leaf B supplies $s = 3, 4$. (b) Representative full cavity trajectories acquire every root. (c) Of 42 initial-power and polarity combinations, 33 acquire a root, four reach the lower imposed power limit, and five reach the upper limit. None remain unlocked. (d) The predicted local rates vary by more than an order of magnitude, and the slowest-conditioned root carries the largest finite-time error. Repeating the four representative runs at time steps 0.01, 0.005, and 0.0025 leaves every classification unchanged. The maximum final-power shift between the two finest steps is 3.16×10^{-12}.

No fixed polarity in Eq. (8) can be restoring from both directions. One side approaches the root while the other is driven towards an imposed power limit. Figure 7 shows the susceptibility collapse, the slowing simple root acquisitions, and all four side-polarity combinations at repetition.

The product hierarchy introduces a second conditioning channel. At a root contributed by leaf A, differentiation of Eq. (10) gives

$$\partial_s E|_{e_A=0} = \frac{\partial_s e_A}{\epsilon_A} \frac{e_B}{\epsilon_B}\Big|_{e_A=0},$$

with the roles exchanged for a root of B. The outer rate therefore depends on the active leaf slope and on the off-root value of the inactive leaf. Normalisation can balance their typical magnitudes, but it cannot recreate a derivative that has vanished. This explains the order-of-magnitude spread in the four predicted rates of Fig. 6. The derivation and integration details are collected in Appendix B.

These limitations are computational, not incidental engineering defects. Delay restricts how aggressively the residual can be followed [24, 25], finite power restricts which roots can be represented, polarity selects a local basin, and multiplicity can remove two-sided accessibility altogether. An optical root is consequently defined by both an equation and an operating domain. The field finds the root only where the material responses continue to provide directional information.

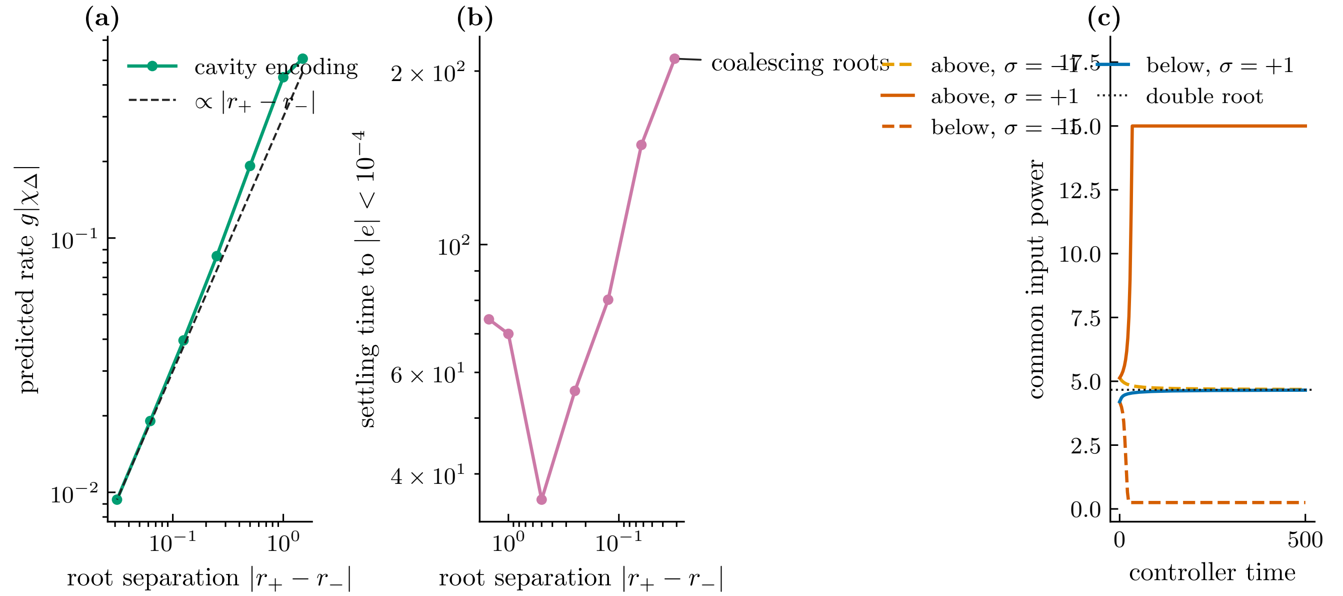

FIG. 7. The conditioning calculation shows the loss of first-order information as two encoded roots coalesce. (a) The predicted local feedback rate vanishes linearly with root separation. (b) All seven simple root cases settle, but the measured settling time grows from 36.09 to 210.14 over the tested separations. (c) At exact repetition each fixed polarity approaches the root from only one side. The opposite-side trajectory reaches the imposed input power limit. The obstruction belongs to the scalar differential response rather than to the choice of gain.

VIII TWINNING AS COMPUTATION

A twinning field is a zero of measured differential response. Section II established the fixed-point equivalence for a strict common-field loop, and Sec. III showed that local acquisition is governed by differential susceptibility and delay. The same mechanism locked the currents of two small Hubbard systems in Sec. IV. Sections V and VI then reversed the original question. Instead of calculating a field for a given pair of materials, we chose the response laws so that their twinning fields represented the roots of a supplied equation. The resulting three-cavity, seven-variable model acquired all four roots of a factorised quartic. Optical indistinguishability is no longer only the target of the control: it is the operation that computes.

The construction looks familiar from several directions because none of its neighbouring languages is ornamental. At the detector, it is integral feedback acting on an optical error signal [8, 9]. As a numerical method, it is a continuous-time root finder acting on a residual. In analogue photonics, it combines nonlinear response elements with a multiplier hierarchy. Material operators already perform mathematical transformations [10] and solve integral equations [11, 14]; optical feedback resonators solve differential equations [12, 13]. Delay reservoirs compute with recurrent physical dynamics [15, 16], while optoelectronic and analogue optical machines seek optimisation fixed points [17, 18]. Programmable or layered optics already supply polynomial transformations [21, 23]. The novelty lies in what evaluates the residual: equality between distinct systems under the same adaptive field.

There is a small trap in the title. The tree does not search its branches. Every leaf is driven at once, each evaluates its own material equality under the same s, and the product of the returned differences determines the next change in s. Calling the construction a twinning tree records this physical organisation and nothing more. We demonstrated its two-leaf form here. A deeper tree would not be this figure with more foliage: in our construction it would raise new dynamic-range, noise, normalisation, and conditioning questions. Those are predictions of the architecture, not results demonstrated here.

The distinction between feasibility and advantage is especially important. Signed optical detection and response-based feedback have direct precedents [6–9]. Driven nonlinear resonators are established physical systems [26, 27, 29], as are programmable nonlinear response and analogue optical multiplication [22, 23]. These ingredients make a laboratory implementation plausible. They do not establish experimental noise tolerance, useful scaling, or a speed advantage. The demonstrated computational utility is structural: supplied coefficients select physical response laws, the systems evaluate the residual continuously, and the retained common field identifies an accessible solution.

Section VII also shows why the physical embodiment cannot be separated from the calculation it performs. Root separation determines susceptibility, the other leaves rescale the local slope, delay limits usable gain [24, 25], and finite input power truncates the search domain. A repeated root removes the sign change needed for two-sided fixed-polarity acquisition; an independently unstable cavity branch could remove the operating point before the polynomial does [28]. These are not reasons to abandon the construction. They are the quantities that any larger one must control.

Several routes remain open. A deeper hierarchy could represent a larger supplied factorisation, but its resource and conditioning scaling must be calculated rather than inferred from the schematic. A slowly deformed family of twinning conditions might permit homotopy-style continuation between roots, although no such controller has been tested. Arbitrary expanded-polynomial encoding, automatic factorisation, complex roots, and multimode response Jacobians likewise require new physical mappings. The original twinning-field problem asked whether one drive could make two materials answer alike [5]. Work on optical and Hamiltonian-encoded reservoirs asks what a physical response can calculate [19, 20]. Here those questions meet: equality between the answers is the computation, and the common field plays three roles at once—it interrogates the materials, carries the iteration, and remains as the returned root.

Appendix A: The scalar loop and its delay boundary

We expand the differential response of Eq. (1) about a regular root as $e(\Phi_\star + \delta\Phi) = \chi_\Delta \delta\Phi + O(\delta\Phi^2)$. At $\tau = 0$, the polarity choice $\sigma\chi_\Delta > 0$ in Eq. (2) yields $\delta\dot{\Phi} = -\kappa|\chi_\Delta|\delta\Phi$. The linear perturbation therefore contracts as $\delta\Phi(t) = \delta\Phi(0)\exp[-\kappa|\chi_\Delta|t]$.

For $\tau > 0$, we substitute $\lambda = i\omega$ into the characteristic equation in Eq. (4). Separating real and imaginary parts gives $\kappa|\chi_\Delta|\cos(\omega\tau) = 0$ and $\omega - \kappa|\chi_\Delta|\sin(\omega\tau) = 0$. The first positive crossing has $\omega\tau = \pi/2$ and $\omega = \kappa|\chi_\Delta|$, hence $\kappa|\chi_\Delta|\tau = \pi/2$. The branch connected continuously to the stable root at $\tau = 0$ is stable below this first crossing.

For the numerical sweep, we obtain the asymptotic exponent from a least-squares fit to $\log|e(t)|$ on the retained tail interval. Negative fitted exponent denotes contraction. We classify finite-time acquisition only when the residual remains below its tolerance for the prescribed hold interval. Sweep coordinates are reported as the dimensionless product $\kappa|\chi_\Delta|\tau$, allowing direct comparison with Eq. (4).

Appendix B: Cavity mappings and numerical protocols

For the Hubbard calculation, each four-site ring is restricted to two spin-up and two spin-down fermions, giving $\binom{4}{2}^2 = 36$ basis states. The Hamiltonian is Eq. (5), and the current is obtained by differentiating it with respect to the common Peierls phase. The two states are propagated independently while the delayed current difference updates the one phase through Eq. (2). Initial states are prepared separately for the two interactions. Lock is assessed on the retained tail by the root-mean-square current mismatch divided by the root-mean-square common current. The instantaneous comparison root is obtained by an independent scalar inversion at each sampled state, not by reusing the feedback trajectory.

For the single-quadratic construction, substituting the parameters of Eq. (7) into the two laws in Eq. (6) gives $P_1 - P_2 = n[(\Delta_1 - n)^2 + \gamma_1^2 - \gamma_2^2] = n(n^2 + bn + c)$. In our parameterisation the linewidths are positive when $L > 0$ and $L + \Delta_1^2 - c > 0$. The working interval is also chosen so that both $P_j'(n) > 0$, making $n_j(s)$ single valued there. The full complex amplitudes, rather than adiabatically eliminated intensity curves, are integrated in the acquisition runs governed by Eq. (8).

For a supplied factor $q_j(s) = s^2 + b_j s + c_j$, substitute the detuning and linewidth in Eq. (9) into $P_j - P_0$. Collecting powers of n gives $nK_j^2 q_j(Gn)/G^2$. At a root contributed by leaf A, differentiation of Eq. (10) gives

$$\partial_s E|_{e_A=0} = \left.\frac{\partial_s e_A}{\epsilon_A}\frac{e_B}{\epsilon_B}\right|_{e_A=0},$$

with the analogous expression after exchanging A and B. The active-leaf slope and the off-root response of the other leaf therefore jointly determine conditioning, consistent with Eq. (11).

All numerical producers write tables before any plotting step. The scalar delay, Hubbard, cavity, conditioning, and twinning tree projectors read retained CSV files and do not invoke the simulations. The two-leaf cavity model has three complex amplitudes and one real controller state, hence seven real dynamical variables. The twinning tree sweep contains 42 initial-power and polarity combinations, uses imposed lower and upper input power limits, and classifies acquisition by retained root-distance and hold tolerances. Representative trajectories were repeated at steps 0.01, 0.005, and 0.0025. Producer hashes, integration horizons, convergence summaries, reference roots, run classifications, and plot-ready traces are retained with the numerical data.

ACKNOWLEDGMENTS

This paper was developed with Eigen through the EigenEngine pipeline. A front-to-back record of the generation, including the conceptric, numerical evidence, manuscript, and companion talk, is available as an EigenEngine featured generation.

[1] G. McCaul, C. Orthodoxou, K. Jacobs, G. H. Booth, and D. I. Bondar, Driven imposters: Controlling expectations in many-body systems, Physical Review Letters **124**, 183201 (2020), arXiv:1911.05006.

[2] G. McCaul, C. Orthodoxou, K. Jacobs, G. H. Booth, and D. I. Bondar, Controlling arbitrary observables in correlated many-body systems, Physical Review A 10.1103/PhysRevA.101.053408 (2020), arXiv:1912.06173.

[3] A. B. Magann, G. McCaul, H. A. Rabitz, and D. I. Bondar, Sequential optical response suppression for chemical mixture characterization, Quantum **6**, 626 (2022), arXiv:2010.13859.

[4] G. McCaul, A. F. King, and D. I. Bondar, Non-uniqueness of non-linear optical response, Annalen der Physik 10.1002/andp.202100523 (2022), arXiv:2110.06189.

[5] G. McCaul, A. F. King, and D. I. Bondar, Optical indistinguishability via twinning fields, Physical Review Letters 10.1103/PhysRevLett.127.113201 (2021), arXiv:2103.01162.

[6] J. Masur, D. I. Bondar, and G. McCaul, Dynamical generation of epsilon-near-zero behaviour via tracking and feedback control, arXiv preprint 10.48550/arXiv.2301.12069 (2023), arXiv:2301.12069.

[7] V. Bilokon, E. Bilokon, and D. I. Bondar, All you need is amplifier: Spectral imposters without pulse shaping, arXiv preprint 10.48550/arXiv.2603.05417 (2026), arXiv:2603.05417.

[8] T. W. Hänsch and B. Couillaud, Laser frequency stabilization by polarization spectroscopy of a reflecting reference cavity, Optics Communications **35**, 441 (1980).

[9] E. D. Black, An introduction to pound–drever–hall laser frequency stabilization, American Journal of Physics **69**, 79 (2001).

[10] A. Silva, F. Monticone, G. Castaldi, V. Galdi, A. Alù, and N. Engheta, Performing mathematical operations with metamaterials, Science **343**, 160 (2014).

[11] N. M. Estakhri, B. Edwards, and N. Engheta, Inverse-designed metastructures that solve equations, Science **363**, 1333 (2019).

[12] J. Cederquist and S. H. Lee, Coherent optical feedback for the analog solution of partial differential equations, Journal of the Optical Society of America **70**, 944 (1980).

[13] J. Hou, J. Dong, and X. Zhang, Optical solver for a system of ordinary differential equations based on an external feedback assisted microring resonator, Optics Letters **42**, 2310 (2017).

[14] A. Cordaro, B. Edwards, V. Nikkhah, A. Alù, N. Engheta, and A. Polman, Solving integral equations in free space with inverse-designed ultrathin optical metagratings, Nature Nanotechnology **18**, 365 (2023), arXiv:2202.05219.

[15] L. Appeltant, M. C. Soriano, G. Van der Sande, J. Danckaert, S. Massar, J. Dambre, B. Schrauwen, C. R. Mirasso, and I. Fischer, Information processing using a single dynamical node as complex system, Nature Communications **2**, 468 (2011).

[16] K. Vandoorne, P. Mechet, T. Van Vaerenbergh, M. Fiers, G. Morthier, D. Verstraeten, B. Schrauwen, J. Dambre, and P. Bienstman, Experimental demonstration of reservoir computing on a silicon photonics chip, Nature Communications **5**, 3541 (2014).

[17] F. Böhm, G. Verschaffelt, and G. Van der Sande, A poor man's coherent ising machine based on opto-electronic feedback systems for solving optimization problems, Nature Communications **10**, 3538 (2019).

[18] K. P. Kalinin, J. Gladrow, J. Chu, *et al.*, Analog optical computer for AI inference and combinatorial optimization, Nature **645**, 354 (2025).

[19] G. McCaul, K. Jacobs, and D. I. Bondar, Towards single atom computing via high harmonic generation, The European Physical Journal Plus 10.1140/epjp/s13360-023-03649-3 (2023), arXiv:2104.06322.

[20] G. McCaul, J. S. Totero Gongora, W. Otieno, S. Savel'ev, A. Zagoskin, and A. G. Balanov, Minimal quantum reservoirs with hamiltonian encoding, Chaos: An Interdisciplinary Journal of Nonlinear Science 10.1063/5.0282921 (2025), arXiv:2505.22575.

[21] M. Yildirim, N. U. Dinc, I. Oguz, D. Psaltis, and C. Moser, Nonlinear processing with linear optics, Nature Photonics **18**, 1076 (2024).

[22] T. Fujita, Y. Toba, Y. Miyoshi, and M. Ohashi, Optical analog multiplier based on phase sensitive amplification, in *2013 18th OptoElectronics and Communications Conference Held Jointly with 2013 International Conference on Photonics in Switching* (Optica Publishing Group, 2013) p. TuPS_8.

[23] T. Wu, Y. Li, L. Ge, and L. Feng, Field-programmable photonic nonlinearity, Nature Photonics **19**, 725 (2025).

[24] K. Pyragas, Continuous control of chaos by self-controlling feedback, Physics Letters A **170**, 421 (1992).

[25] N. D. Hayes, Roots of the transcendental equation associated with a certain difference-differential equation, Journal of the London Mathematical Society **s1-25**, 226 (1950).

[26] H. M. Gibbs, S. L. McCall, and T. N. C. Venkatesan, Differential gain and bistability using a sodium-filled fabry-perot interferometer, Physical Review Letters **36**, 1135 (1976).

[27] P. D. Drummond and D. F. Walls, Quantum theory of optical bistability. i. nonlinear polarisability model, Journal of Physics A: Mathematical and General **13**, 725 (1980).

[28] K. Ikeda, Multiple-valued stationary state and its instability of the transmitted light by a ring cavity system, Optics Communications **30**, 257 (1979).

[29] L. A. Lugiato and R. Lefever, Spatial dissipative structures in passive optical systems, Physical Review Letters **58**, 2209 (1987).